DOCUMENT RESUME

ED 328 483 SO 021 067

AUTHOR Baker, Mike
TITLE Intelligent Computer-Aided Instruction and Musical Performance Skills. CITE Report No. 18.
INSTITUTION Open Univ., Walton, Bletchley, Bucks (England). Inst. of Educational Technology.
PUB DATE Jul 87
NOTE 26p.; Transcript of a talk presented at the Conference of the Open University CAL Research Group (Milton Keyes, England, December 1986). For related documents, see SO 021 064-066.
PUB TYPE Reports - Descriptive (141) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Cognitive Processes; *Computer Assisted Instruction; Computer Uses in Education; Educational Technology; Foreign Countries; Interactive Video; Learning Strategies; Learning Theories; *Music Education; Music Techniques; *Music Theory; Skill Development; Teaching Methods

ABSTRACT

This paper is a transcription from memory of a short talk that used overhead projector slides, with musical examples played on an Apple Macintosh computer and a Yamaha CX5 synthesizer. The slides appear in the text as reduced "icons" at the point where they would have been used in the talk. The paper concerns ways in which artificial intelligence techniques could be applied to enable computers to make some kind of useful input into teaching musical performance skills. The overall plan of the paper falls into three main parts, namely: "What should be taught?" "How could it be taught?" and "Could it and should it be done?" A particular psychological characterization of musical performance is discussed as a means of answering the first question: what should teachers be trying to teach? The question, how could teachers teach it, then becomes a problem of finding a knowledge representation which both models these cognitive processes and is suitable for the teaching strategies desired. The question as to whether the proposal could be done is a matter of having the needed human and technological resources--and of stretching them to the limit in anticipation of future advances. As to whether it should be done, what is needed is a reasonable assurance that the research will make some useful contribution both to music education and to general design principles in compouter aided instruction. Design considerations for the technology for creating the appropriate environment to develcp music performance skills are examined and practical implications of the human and technological resources used for computer design and for music education are explored. An appendix contains questions, suggestions, and answers that were raised in the original discussion, together with a short list of works referred to in the paper. (DB)

SCOPE OF INTEREST NOTICE

The ERIC Facility has assigned this document for processing to: SO

In our judgment, this document is also of interest to the Clearinghouses noted to the right. Indexing should reflect their special points of view. IR

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
C. J.
LAWLESS
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

SO 021 067

INTELLIGENT COMPUTER-AIDED INSTRUCTION AND MUSICAL PERFORMANCE SKILLS

Introduction

What follows is a transcription from memory of a short 20 minute talk given at the Open University CAL Research Group Conference, December 1986. The talk was given using over-head projector slides, with musical examples played on an Apple Macintosh computer and Yamaha CX5 synthesiser, using commercially available composition and performance software. It was assumed that members of the audience would have knowledge of use of artificial intelligence techniques in education, but that they would have little detailed knowledge of music theory and performance. Since the ideas contained in the talk were culled from a much longer Open University (C.I.T.E.) technical report ('Design Considerations for an Intelligent Tutoring System for Musical Performance Skills'), it was felt that in the interests of brevity it would be more useful for the reader to be presented with the visual materials used in the talk together with an attempt at a faithful verbal reconstruction. Accordingly, slides used appear within the text as reduced 'icons', at the point where they would have been used in the talk. An abstract is not provided since a quick overview can be obtained by scanning the icons.

As an appendix, a few questions, suggestions and purported answers which were raised in discussion are included, together with a short list of works referred to.

Due apologies (should the reader require them) are made for use of informal conversational English throughout. Finally, I must express my warm thanks to the many people in the Institute of Educational Technology who had the patience to listen to my many earlier atttempts at giving this talk, and for their constructive criticism. Special thanks to my supervisor and co-supervisor - Mark Elsom-Cook and Ann Jones - to Simon Holland, Pat Fung and Sarah Hennessy. All responsibility for errors, inadequacies and misunderstandings lies with myself.

Structure

Written documents can encourage a linear interpretation of the ideas presented, which are often of a branching structure in a talk, since many slides are used and re-used as further slides enlarge upon a particular area. For this reason, in view of the somewhat unusual format of this document, the following diagram may be of use in understanding the branching structure of the talk.

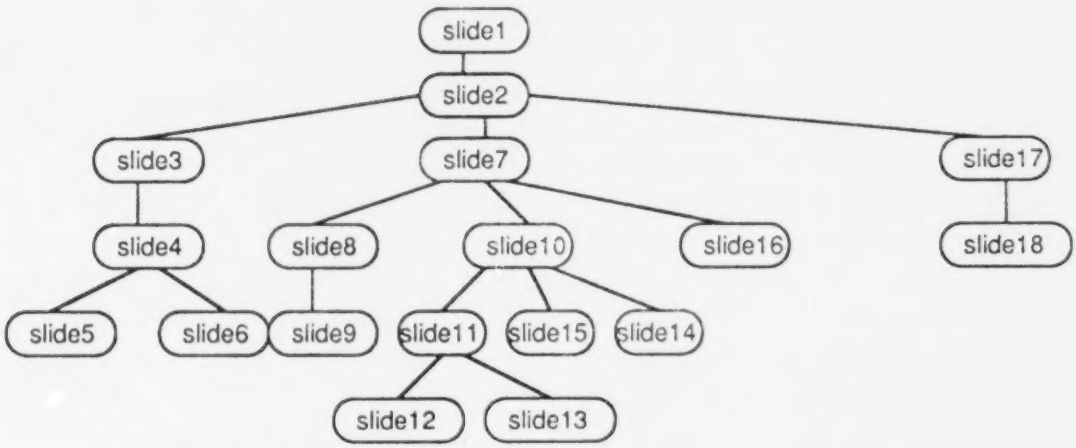

SLIDE 1

TUTORING MUSICAL PERFORMANCE SKILLS

*OBJECTIVES

'To determine how computers and artificial intelligence techniques can assist learning of musical performance skills by novice instrumentalists'

*KEYWORDS

- COMPUTERS

- AI

-MUSICAL PERFORMANCE SKILLS ('MPS')

-NOVICES

My present research objectives are to examine ways in which we could apply artificial intelligence techniques to enable computers to make some kind of useful input into teaching musical performance skills. The general idea is primarily to help *novices* , which is one of the reasons why we might want to use computers in the first place, perhaps by allowing the machine to take over some activities for the student whilst leaving him free to concentrate on others. By saying this I do not want to exclude the possibility that experts could derive some assistance from the projected educational tools and devices, given that conventional (MIDI) instruments could be interfaced to a tutoring system.

Presently I see AI techniques as providing a powerful deep level knowledge representation for music in the context of performance and for modelling the human cognitive processes involved in musical performance. At this point I draw attention to the key subject area of musical performance skills ('MPS' for short) since one of my preliminary objectives is to determine exactly what psychological processes are involved in this diffuse group of skills, and hence which of them we should in fact be teaching: clearly a tutoring system should not simply aim to reduplicate present conventional ways of teaching without attempting to determine their educational worth and should attempt to teach those areas which are most suited to the unique possibilities offered by computer-based systems.

SLIDE 2

A GUIDED DISCOVERY ENVIRONMENT
FOR MUSICAL PERFORMANCE SKILLS

PLAN:

(1) WHAT ARE WE TRYING TO TEACH?
-MUSIC EDUCATION
- A MUSICAL EXAMPLE
- PSYCHOLOGICAL EVIDENCE

(2) DESIGN CONSIDERATIONS FOR
A GUIDED DISCOVERY ENVIRONMENT
FOR MPS
- KNOWLEDGE REPRESENTATION
-TEACHING STRATEGIES

(3) CONCLUSIONS
- WHAT'S IN IT FOR ITS DESIGN
- WHAT'S IN IT FOR MUSIC EDUCATION
-PRACTICAL POSSIBILITIES

The overall plan of what I want to say falls into three main parts, which may be expressed as

'what should we be teaching?'

'how could we teach it?'

'could it and should it be done?'

In order to decide what to teach I shall consider a particular psychological characterisation of musical performance (although a detailed consideration of suporting evidence and theoretical jusification is beyond the scope of this talk).

The question as to *how* to teach then becomes a problem of finding a knowledge representation which both models these cognitive processes and is suitable for the teaching strategies which we would like to use it for.
As to whether the proposal *could* be done, this is a matter of the human and technological resources which are at my disposal, although a research project should nevertheless stretch present technologies to their limits in anticipation of future advances. Lastly, we need to have reasonable assurance that research will make some useful contribution to both music education and to general design principles in intelligent computer-aided instruction: we do not want to actually make things worse for a student of musical performance by producing an efficient implementation of a bad educational idea, and neither do we want to produce yet another ad hoc programming solution to another subject-area implementation.

SLIDE 3

WHAT ARE WE TRYING TO TEACH?

*TRADITIONAL EDUCATIONAL APPROACHES:
- THEORY
- NOTATION
- SCALES AND ARPEGGIOS
-AURAL TRAINING
-PERFORMANCE: TECHNIQUE
INTERPRETATION

*CONCLUSION: A LOT OF WHAT TAKES PLACE UNDER THE HEADING OF MPS IS CONCERNED WITH SOMETHING OTHER THAN MUSICAL PERFORMANCE

*SUPPORTIVE AND INTRINSIC SKILLS
-THEORY: need a vocabulary
-NOTATION: intrinsic?
-SCALES & ARPEGGIOS : technical
-AURAL TRAINING: which kind?
-PERFORMANCE:
TECHNIQUE - some connection between expression and physical gesture
INTERPRETATION - how can we characterise this?

A likely place to look for what we should be trying to teach would seem at first sight to be in the syllabuses of traditional exams in musical performance (such as the 'grade' examinations set by the Associated Board for various instruments). These include theory, sight-reading, scales and aural training, as well as the activity of performance itself, which may be rather simplistically divided into technique (actual production of the right notes) and interpretation (the art of shaping the notes to give a musical effect).

The general conclusion to which we are driven is that a lot of what takes place under the heading of musical performance skills is concerned with something other than performance itself, so the natural question arises as to what the educational purpose of these other activities is intended to be? Clearly, they are intended to be in some way *supportive* to the central activity of performance. This, however is a contentious issue. Although there is not time to discuss current trends in music education and relevant psychological evidence in detail, overall the following view is suggested.

There does seem to be a current trend in music education in this country which is moving away from academic musical activities which are of questionable relevance to musical practise, towards concentrating on the activities of performance and composition themselves (such as the current GCSE syllabus). In fairness to traditional 'back to basics' educationalists, activities such as identifying (i.e.*naming*) modulations and intervals may be of more relevance to the aspiring composer. For the most part I would say that most of these additional activities are concerned with *facility with conventional musical notation*, which may not be the most suitable musical representation for a computer-based tutoring system aimed at novices. It is more useful to the performer to develop musical memory (mentally hearing a note a fraction before you play it) rather than developing an ability to assign conventional names, and a special form of aural training which may be termed 'advanced musical perception', involving the

ability to make fine aural discriminations in terms of tempo, timbre etc. between different performances.

If we turn to the actual activity of musical performance we find that we cannot entirely divorce technique from interpretation since there is clearly some connection between physical gesture and musical expression.

For example, production of notes with increasing pitch usually requires more physical effort, especially with the voice or brass instruments. Commonly this means that it is easier to produce higher notes if they are played more loudly, and this fact has become incorporated into generally accepted musical practise in Western music.

Turning to interpretation of music our problem now is to find a way of characterising this.

SLIDE 4

LINGUISTIC STRUCTURAL ANALYSES AND PERFORMANCE RULES

*STRUCTURAL ANALYSES
(LERDAHL AND JACKENDOFF ,1983)
- LINGUISTIC ANALYSES OF MUSICAL INTUITIONS
-GENERATIVE GRAMMER OF POSSIBLE STRUCTURAL DESCRIPTIONS
- HIERARCHICAL
- STRUCTURAL INDETERMINACY
- PERFORMER COMMUNICATES A PARTICULAR STRUCTURE

*PERFORMANCE RULES
(SUNDBERG, 1984)
- HYPOTHESISED FROM MUSICIANS' INTUITIONS
- RELATED TO
(1) MODELLING INSTRUMENTS
(2) COMMUNICATING STRUCTURE ('GESTALT'CHUNKING')
(3) CULTURAL EXPECTATIONS (eg FINAL RETARD)

MJB 1986

I propose to describe without preamble one particular psycho-linguistic way of characterising performance which appears to be a promising line of research. Fundamentally, it is the view that we can consider music to be analogous to a natural language, and hence that we can see performance as a matter of acquiring the grammatical structure of a piece and being able to reproduce this in performance by applying a set of *rules* for musical expression which operate on this abstract cognitive structural representation of the music.

However simple the model appears in its broad outline, the detailed working of these theories is far from simple! Most recent work in what may be termed 'musico-linguistics' takes as its starting point the work of Lerdahl and Jackendoff in their book A Generative Grammar of Tonal Music, 1983 (see references), which is a kind of amalgamation of Schenker's ideas on structural analysis of tonal music with Chomskyan ideas on language and thinking. Thus Lerdahl and Jackendoff see their theory as providing a linguistic analysis of tonal music which accords with the musical intuitions of some kind of idealised listener. In Chomskyan terms, they seek to provide a 'generative grammar' which will generate possible structural descriptions for pieces of music rather than actually generating pieces of music. I understand from conversation with my colleague Simon Holland that they have now relaxed this disclaimer somewhat to accept that their theory could be used for musical synthesis. Structural descriptions, in their theory do not however work simply on one level - they are hierarchical, on the levels of melody, harmony and rhythm, and structural divisions as specified on one level may sometimes conflict with those predicted on another i.e. their theory admits *structural indeterminacy*. (Incidentally, when I talk of structual divisions in music I'm talkly broadly about divisions such as phrases, and smaller divisions which are called 'motifs' - think of the famous opening motif of Beethoven's fifth symphony, which pervades the whole piece.) For example, two chords which are termed dominant and tonic usually signal the end of a fairly significant structural division, such as a phrase; but given a certain metric occurrence, they do not - this happens often in Bach' chorale harmonisations.

So where, you might ask, does the performer come in? The answer is quite closely related to the indeterminacy which I have discussed: although Lerdahl and Jackendoff state that their concern is not with the role of the performer, they nevertheless state that the performer can determine *which* of possible structures a listener actually hears by placement of subtle pauses and dynamic and other variations. To put it another way, one way of characterising the function of the performer is to *communicate musical structure*. This provides the vital link between this theory and musical performance. Let us consider how a performer could actually do this. The most recent studies of performance have been performed by Sundberg (see references) and co-workers at the Royal Swedish Institute of Technology. Their approach involves formalising rules for variation of such factors as tempo and dynamics hypothesised by an expert performer, and applying these rules using a computer to pieces of music in an attempt to improve the musicality of the result - thet term it 'analysis by re-synthesis'. Psychological experiments are then performed to test the results by playing the re-synthesised pieces to music students and asking them to perform an evaluation. In summary, they found that as well as there being rules which assisted in communicating structure to listeners (which would appear to tie in quite nicely with Lerdahl and Jackendoff's theories) there are other rules which are concerned with modelling traditional instruments - especially the voice, which appears to be the root of a rule which says 'the higher the louder' - and rules which fulfill culturally determined expectations - such as the final retard. Sundberg's work forms an ongoing research project to study the psychological processes involved in performance.

Let's look now briefly at a short musical extract designed to give an outline of how rules can be applied in relation to structure.

SLIDE 5

RULES APPLIED TO STRUCTURE

(From Mozart Piano Sonata K. 331, quoted in Lerdahl & Jackendoff, 'A Generative Theory of Music' (1983))

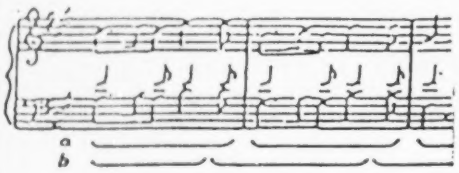

(* = rule experimentally tested by Sundberg et al)

R1: if <abrupt musical parameter change at structural boundary which implies structural analysis S1>
then <apply S1 consistently>

R2*: if <phrase-end>
then <insert micropause>

R3: if <texture = 'melody/accomp'>
then <melody louder-than accomp>

R4: if <texture = greater-than one voice>
then <apply synchronisation>

The example comes from a well-known Mozart piano sonata, and is one which Lerdahl and Jackendoff cite as an example of structural indeterminacy. First I'll play a literal realisiation of the notation, which is what most commercially available software would do, and which is musically unintelligible! Such software even contains a 'quantization' facility for getting rid of all those unwanted performance variations. If you look at the music here you'll see that there are two plausible ways of chunking the extract, labelled 'a' and 'b'. In terms of symmetry Lerdahl and Jackendoff prefer 'a' , but strangely enough most people I've played the extract to prefer 'b' and this is what most performers seem to do. We can influence your perception of grouping quite simply by inserting a micropause - about a demi-semiquaver - and shortening the last note of the group correspondingly. If you listen to the two extracts where I've inserted the pauses at the right points, it seems to me quite surprising how such a small alteration can radically alter your perception of grouping. Inserting the micropause corresponds to what I've termed rule 2, and rule 1 is necessary to secure consistency in structural interpretation. In addition I've included two other rules which could be applied (there are many other possibilities).

We now have a broad characterisiation of musical performance which has psychological justification , and which we can hopefully use to help us decide what to teach.

SLIDE 6

WHAT WE SHOULD BE TEACHING

*ACQUISITION OF PERFORMANCE PLANS WHICH ARE POSSIBLE STRUCTURAL DESCRIPTIONS OF THE MUSIC

*RULES FOR EXPRESSIVE VARIATION OF PERFORMANCE PLANS WHICH ARE WITHIN THE SPACE OF POSSIBLE INTERPRETATIONS

*AURAL ABILITY TO DISTINGUISH DIFFERENT INTERPRETATIONS

MJB 1986

In summary to the first section of my talk, the following three things should be taught by a computer-based tutoring system for musical performance skills:

(1) acquisition of an appropriate conception of the structure of the piece to be played (a 'performance plan') , and generally, ways of psychologically 'chunking' the musical language which are musically acceptable to listeners;
(2) rules (not necessarily known explicitly as such) for variation of the performance plan which improve the musicality of the result and which are within what may be termed the 'space of possible interpretations';
(3) aural ability to distinguish different interpretations.

SLIDE 7

DESIGN CONSIDERATIONS FOR
A GUIDED DISCOVERY ENVIRONMENT
FOR INTERPRETATION OF MUSIC

*KNOWLEDGE REPRESENTATION

- AIMS
-WHAT TYPE OF DOMAIN?
-REPRESENTING PERFORMANCE RULES
-TYPES OF RULES
-RULE MATCHING
-CONFLICT RESOLUTION
-RULE FIRING

*TEACHING STRATEGIES

-POSSIBLE STRATEGIES
-INTERFACES
-RELATION TO EXISTING SYSTEMS
-ELABORATIONS ON STRATEGIES

This slide is simply a plan of the topics to be covered in my second section. It falls into two sections, concerned with finding a computational representation for musical structures and performance rules, and with strategies for using this representation to teach the objectives stated in the first section. In the first subsection I examine what is basically a production system approach, and try to assess how well this model fits the subject domain of musical performance. In the second subsection, I examine an overall teaching approach which is specifically designed for our objectives, and discuss its relationships with existing tools and teaching systems in the field of musical performance skills.

SLIDE 8

KNOWLEDGE REPRESENTATION

*AIMS
-PSYCHOLOGICAL VALIDITY
-SUPORTS TUTORIAL STRATEGIES
*WHAT TYPE OF DOMAIN?
-OPEN: DIFFERENT ANALYSES
DIFFERENT RULES
*REPRESENTING PERFORMANCE RULES
-FRIBERG'S : 'RULLE'
*TYPES OF RULES
-SURFACE & CONCEPTUAL LEVELS
-PSYCHOLOGICAL EXPLANATION
*RULE MATCHING
-WHICH RULES APPLY?
- RULES STUDENT HAS APPLIED
- AUTOMATED ANALYSIS?
- PLURALITY OF REPRESENTATIONS
(CF EBCLIOGLU)
*CONFLICT RESOLUTION
-DON'T NEED GLASS BOX EXPERT
- EXPLORATION OF CONTEXT
(ELABORATIONS - SIMULATIONS)
*RULE FIRING
-NOT FOR EXPERT SOLUTION
(IDEAL BY CO-OPERATIVE EFFORT)
-EXAMPLE GENERATION

With any tutoring system you might expect the form of knowledge representation to truly represent the domain in question - in this case, given a certain conception of artificial intelligence, to model human cognitive processes - and to be suitable for the purposes for which it was intended - i.e. to support proposed tutorial strategies. Since both proposed representation
and tutorial strategies are conjectural at this point,I feel willing to modify both should insuperable difficulties be found. Ideally, of course, our representation should be infinitely flexible, completely separable from teaching strategies, and able to support *any* teaching strategy!

The domain of musical performance present a number of problems which would be faced by any attempt to use the representation to determine single expert solutions as to how pieces should be performed : the domain is 'open' in two senses, in that there are a number of possible structural analyses for most pieces, and furthermore, there are many different possible rules which could be applied. Thinking back to our brief example from Mozart's piano music, there are many different types of rules. A surface level categorisation given by Sundberg was given (slide 4) in terms such as instrumental modelling. However, if we are to use rules to teach we need deeper level categorisations and representations on both *conceptual* and *psychological* levels. There are lessons to be learned from past intelligent tutoring systems here: experience with 'GUIDON' based on an expert system shell showed that students do not find explanations in terms of pointing to seemingly ad hoc rules educationally illuminating,

and the exhaustive list of errorful mathematical problem solving rules of 'BUGGY' may be reduced to a smaller number of underlying misconceived concepts. In the area of performance rules, such underlying explanatory concepts could be 'consistency', 'coherence' , 'symmetry' or 'variety and unity', the psychological explanation for which surely lies in the Gestalt psychology of pattern perception.

Leaving these thorny problems aside for a moment, a computer-based knowledge representation of Sundberg's performance rules was demonstrated at the International Computer Music Conference this year by Anders Friberg, co-working with Sundberg.

SLIDE 9

'RULLE' (FRIBERG 1986)

*IMPLEMENTATION OF SUNDBERG'S RULES FOR TESTING PSYCHOLOGICAL VALIDITY
*Le_Lisp + MIDI DRIVERS ; APPLE MAC; YAMAHA DX7
* MUSIC REPRESENTATION
- note strings + attached property lists
eg (dr793 nomdr854 a0 f0 bar1 phrase1 key"d")
*RULE REPRESENTATION
(PROPERTY MANIPULATION FUNCTIONS)
RULE CONSTRUCT:

```
(DEFUN<rule name>(<input parameters>)
  (EACH-NOTE-IF
     (<condition 1>)
     (<condition 2>)
     (<condition n>)
           (THEN
                (<statement 1>)
                (<statement 2>)
                (<statement n>) )))
```

-manual editing of property-value pairs
-hierarchical rule structures: triggers subject rules via note properties.
(after Friberg and Sundberg ICMC 1986).

This slide gives some more detailed information on Friberg's implementation, which is not in fact intended as an educational idea, but rather is for the purpose of facilitating experimental rule-application and testing. Nevertheless it illustrates how performance rules can be applied to music using the artificial intelligence programming language Le_Lisp, together with customised utilities for driving synthesisers from this language. Basically, music is represented in the 'RULLE' system as strings of note signs, each with an attached list of properties such as nominal (metrical) duration, actual duration, amplitude etc. Performance rules are then simply functions which manipulate these properties in a rule-like manner. It must be emphasised that in 'RULLE' rules are essentially applied by hand, in that an area of music is selected and a rule applied using familiar Macintosh menu and editing facilities - i.e. a human musician decides the crucial questions as to *which* rules apply, and *where* in the music. As will emerge, it may be too ambitious but neverthless desireable to expect automated structural analysis (possibly of the Lerdahl and Jackendoff kind) in order for a tutoring system to decide where which rule applies (nb. rules generally have structural information to be satisfied in their antecedents).

Returning to our previous slide,assuming this representation, a number of problems would be faced by a production-system approach:

SLIDE 8

KNOWLEDGE REPRESENTATION

- *AIMS
  - -PSYCHOLOGICAL VALIDITY
  - -SUPORTS TUTORIAL STRATEGIES
- *WHAT TYPE OF DOMAIN?
  - -OPEN: DIFFERENT ANALYSES
    DIFFERENT RULES
- *REPRESENTING PERFORMANCE RULES
  - -FRIBERG'S : 'RULLE'
- *TYPES OF RULES
  - -SURFACE & CONCEPTUAL LEVELS
  - -PSYCHOLOGICAL EXPLANATION
- *RULE MATCHING
  - -WHICH RULES APPLY?
  - - RULES STUDENT HAS APPLIED
  - - AUTOMATED ANALYSIS?
  - - PLURALITY OF REPRESENTATIONS
    (CF EBCLIOGLU)
- *CONFLICT RESOLUTION
  - -DON'T NEED GLASS BOX EXPERT
  - - EXPLORATION OF CONTEXT
    (ELABORATIONS - SIMULATIONS)
- *RULE FIRING
  - -NOT FOR EXPERT SOLUTION
    (IDEAL BY CO-OPERATIVE EFFORT)
  - -EXAMPLE GENERATION

Production systems operate in cycles of matching to see which rule(s) applies, conflict resolution between rules, and ultimately firing the rule or rules. For purposes of teaching musical performance, these features of the production cycle would not all be used in a straightforward way. As I mentioned, it would be obviously preferable to determine rule matching by automated structural analysis, but before embarking upon such a major enterprise we need to be sure about its necessity given our proposed teaching strategies. Given a recorded copy of student 'performed' musical input, we would need to be able to determine

which of the possible rules had been applied and to be able to distinguish between correct and incorrect applications. It may turn out that the form of knowledge representation which was suitable for rule application was not equally sutable for such analysis, and that we may require a plurality of representations.
When rules conflict, it may not turn out to be necessary to actually resolve the conflict by providing some kind of preference mechanism , but rather, in view of the nature of performance and our educational purposes, to simply allow the student to explore the context within which the conflict arose. Similarly, we may not need to fire rules in order to produce 'expert solutions' : the openness of the subject domain would seem to exclude the simple approach of comparing student input with a single expert solution. We may however require rule firing for producing many kinds of examples, discussed later in elaborations to basic teaching strategies.
The general conclusion to this subsection is that we need to look closely at our proposed teaching strategies before we leap into elaborate knowledge representation techniques which would far exceed our purposes.

SLIDE 10

POSSIBLE STRATEGIES

* NON-REAL TIME :
PERFORMANCE PLAN ACQUISITION
-GOALS DEFINED BY LEARNER WITH TUTORIAL GUIDANCE
-VISUAL EDITING INTERFACE
- LEARNER-DEFINED CONCEPTS FROM SYSTEM PRIMITIVES

*REAL-TIME:
PERFORMANCE PLANS INTO PRACTISE
- MATCHING OF IDEAL PERFORMANCE WITH REAL TIME PERFORMANCE VIA GESTURAL CONTROL INTERFACE
-TUTORIAL GUIDANCE GIVEN ON BASIS OF MATCH AND ERROR DIAGNOSIS
- NEW MUSICIAN-MUSIC RELATIONSHIP

*COMBINATION OF NON-REALTIME AND REAL-TIME:
-ANALOGY: TAPE RECORDING , & LISTENING TO EXPERTS

3 LEVELS OF INTERACTION:
(1) ANALYSIS
(2) RULE APPLICATION
(3) PERFORMANCE OF IDEAL

I broadly divide possible teaching strategies into non-real-time and real-time systems.
To begin with the former, I want to say that a non-real-time approach corresponds to our first tutorial objective from section 1: *acquisition of correct performance plans.* The general approach advocated is for the student and computer coach to determine tutorial goals by a process of co-operative decision. Decisions to be made concern firstly the way in which the piece of music to be 'performed' in a tutorial interaction can acceptably be analysed into structural units, and how other features of the music (such as harmony and texture) can be identified.

If the student decides upon a way of 'chunking' the piece which is outside the space of possibilities or which is applied inconsistently, the coach will generate appropriate tutorial advice and explanations. Since there are a number of acceptable possibilities, the student may decide on alternatives on the basis of his musical intuitions.

The second stage is to coach the student into applying suitable performance variations to the analysed piece, in an attempt to improve the musical effect of the piece whilst remaining within the space of possibilities as defined by the coach. For instance, the student might decide that a particular unit should have a gradual increase in volume (crescendo). If the student applies a variation which ignores structure at an incorrect point in the music, the coach generates advice, and suggests possibilities which the learner may not have considered and which he is free to accept or reject. By these means the goal of the interaction will be to produce an edit of how a piece could be performed which corresponds to the student's ideal, but which is musically acceptable.

How could this be done in practise? Although analytical and performance choices must be made on the basis of aural intuition, it would nevertheless be useful for the student to have a visual record of the editing process. Thinking of novices, it may be best not to use conventional musical notation, but rather some kind of graphic interface. Over the past decade a number of visual interfaces for computer music have appeared, most of which use some kind of 'piano-roll' notation, where the exact duration of a note is displayed on a conventional stave, allowing visual correction of rhythmic inaccuracies (eg Lamb & Buckley 1982).

An interesting and different approach to interface design was demonstrated by Adrian Freed at this year's International Computer Music Conference, called 'MacMix'.

I'll discuss this in a bit more detail when I come to talk about interfaces in general.

When we come to actually applying performance variations, I'm not envisaging the student actually being presented with rather abstract looking rules as such. The idea is simply to present them with a menu of choices as to things you can do with the music, such as

louder
softer
increasing volume
change tone
etc.

and then promote individualisation by allowing the student to redefine performance concepts (in terms of system primitives which allows rule matching) or by allowing them to 'play' a section with a variation, save it and define it. In this way the student should discover for himself what the rules are by repeated application and misapplication.

Turning now to real-time tutorial strategies, I want to say that this strategy corresponds to our second objective of being able to actually reproduce your ideal by applying performance rules to performance plans in actual real-time performance. Performance is intimately involved with aural memory, so a good way of testing whether the student has actually formed a cognitive performance plan of a piece and can apply performance rules in a principled way, is to get the student to re-perform their ideal edit. We could then generate coaching advice on the basis of matching between real-time performance and ideal edit.

This could lead to such things as diagnosis of causes of errors and generation of appropriate remedial instruction. Performance could obviously be done using any conventional MIDI instrument, but for use with novices a number of 'gestural control' interfaces (discussed in the next slide) have recently been developed which give the musician the kind of control over performance factors which an orchestral conductor would have without actually requiring him to produce the notes.

Clearly then, an ideal learner-discovery coaching environment would combine real-time and non-real-time elements. The best way of thinking of this is by analogy with two pieces of technology which musicians have used since they were invented: the tape recorder and LP recordings. Players often record their performances for purposes of critical analysis, but this has the problem that you can only get so far by tugging on your own bootstraps - eventually your need a teacher in order to help in the critical process. Similarly, musicians often listen to recordings of pieces by expert performers, which can have the disadvantage of encouraging slavish imitation. I would say that the tutorial strategies which I have described incorporate the advantages of both these technological aids, whilst improving on their weaknesses. Students are able to create their own ideal performances for purposes of self imitation, with the advantages that they have coaching guidance in so doing and fully understand the reasons why the piece should be played in that way because they have been through the process of creating that ideal. Furthermore, they can benefit from advice on the reasons why they cannot carry out their ideal.

In summary, then, I am proposing a strategy based on three levels of interaction:
(1) ANALYSIS of musical structures;
(2) APPLICATION of rule-based performance variations;
(3) PERFORMANCE of co-operatively achieved ideal.

SLIDE 11

INTERFACES

*NEW RELATIONSHIP
- INSTEAD OF ALL MUSICAL INFORMATION PASSING THROUGH PERFORMER, COMPUTER AND HUMAN SHARE THIS

*A CRUCIAL ISSUE
- EXTENT TO WHICH PHYSICAL GESTURES ARE ESSENTIAL TO EXPRESSIVE PERFORMANCE

*GRAPHIC
- eg MacMix, piano roll notation (Lamb and Buckley)

*GESTURAL
- SSSP
- SEQUENTIAL DRUM
- INTERACTIVE COMPOSITION
- OCTAPAD

*MIDI
- HARDWARE COMPATIBILITY
- SOFTWARE TOOLS
- PROBLEMS

Firstly, I must point out that with computer music applications (especially those which claim to be educational) choice of interface is crucial. Quite simply, given the connection between physical gesture and expression in performance,

if we use a gestural interface to replace the complex gestures required for playing a conventional instrument, we must make sure that it embodies the right kind of gestures which are involved in musical expression! If we use a visual interface for feedback in real-time or for representing musical concepts in non-real time, we must be sure that visual feedback ais appropriate and that concepts are represented in ways which facilitate individualised learning.
The following is part of the MacMix graphic interface discussed earlier (slide 13). Basically, melodies are represented as amplitude envelopes, and can be edited using facilities similar to many Apple Macintosh word processing packages. Selections from pieces can be made and combined together in pull-down 'mix windows', in combination with other timbres. The parallel processing facilities of computers which we use here such as the Apollo Domain system would have the advantage of permanent menus and windows, minimising the learning time for novice users.
At this stage it is not quite clear exactly what kind of visual interface would be most suitable for my purposes, and I have discussed MacMix as an example of the general type of interface envisaged.

SLIDE 13

MACMIX GRAPHIC INTERFACE

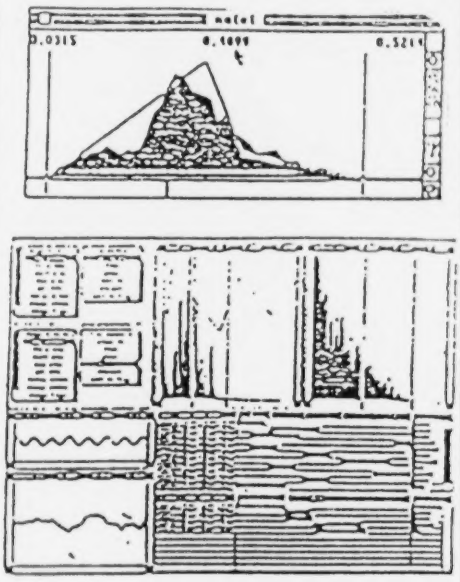

(ADRIAN FREED ICMC 1986)

Turning to gestural control devices, a good example is the 'sequential drum', developed by Mathews and Abbott (1980), in which you can trigger the notes of a complex composition in

real-time by whacking a rubber drum pad which senses velocity and timing of attack, amongst other things.
In this way you can have control over such expressive devices as *tempo rubato* and dynamic variation.
Other examples include the SSSP conducting system (Buxton et al), Chadabe's interactive composition device (where you have to react to unpredictable compositional input and the system senses arm control movements using antennae) and more recently commercially available devices such as the Roland 'octapad' drum pads and Palmtree Instruments 'airdrums' which can control synthesisers via MIDI by moving two hand-held devices like maracas !

The important thing to note about these gestural control devices is that they represent a new kind of relationship between musician, score and instrument, as the next slide shows (slide 12). Traditionally, all information from the score passes through the musician to the instrument, which was superseded in the early phase of computer music by musicians specifying scores as algorithms which the instrument (computer) played. With intelligent instruments, both score (ie the uninterpreted note) and human musician feed information to the computer instrument. Finally, in the proposed discovery environment, both student and coach (together forming the interpretive musician) and the score feed information to the instrument.

SLIDE 12

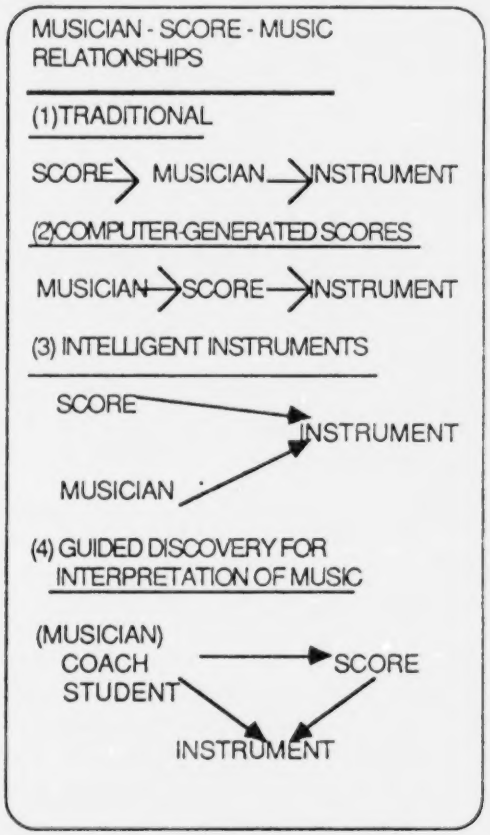

Finally, I must mention the MIDI (Musical Instrument Digital Interface) industry standard for hardware compatibility which has been established in the last few years. Despite alleged limitations in a standard which was not really intended for computer music research requirements, the standard puts real-time performance research within the reach of anyone with a desktop computer and a MIDI synthesiser, and has stimulated the development of software development tools, such as MIDI_Lisp (Wessell, Boynton et al) which enable synthesiser control from an AI programming environment.

SLIDE 14

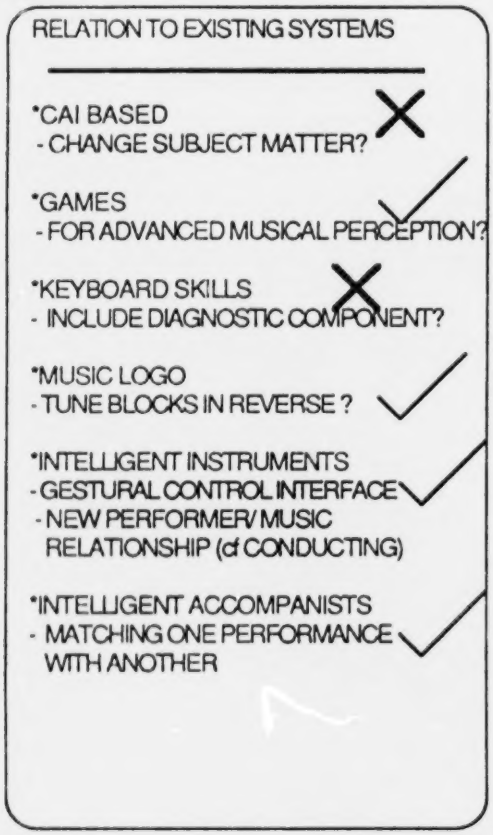

I'd now like to attempt to briefly assess the relation of these proposals to existing work. To begin with what it is not, the proposed system owes little to CAI (computer assisted instruction) based tutoring systems, either those with branching text frames, or those which generate these. The underlying educational approach is rather that of a discovery environment, more akin to LOGO ideas, with an element of coaching. Nor is it influenced by those piano keyboard tuoring systems which view learning as a matter of approximating the student's performance to that of an expert or to a more or less literal rendition of the notation, other than perhaps in utilising similar kinds of interfaces. Turning to closer affinities, the first proposed phase of musical analysis could be likened to the Music LOGO 'tuneblocks' game described by Bamberger in which student recombine familiar nursery rhymes from constituent musical blocks, except that here the reverse activity of specifying the blocks in the first place is advocated.

From aural training games, we could envisage a game with convergent goal definition for advanced musical perception, in which the game is to specify in what respects performances of the same piece differ.

Finally, from intelligent instruments I take the novel conception of the relationship between performer, instrument and score, together with ideas on matching one performance with another from intelligent computer accompanists for live performances (such as the system described by Vercoe, see references).

SLIDE 15

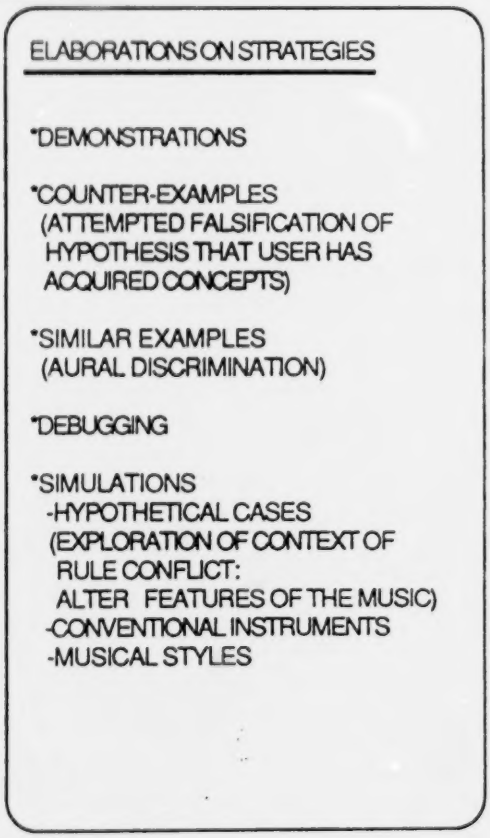

One way in which consideration of existing work may assist my proposal is in suggesting ways of making our tutorial scenarios more concrete in their detail. The following are a few ways of elaborating on these scenarios.

(1) provision of annotated demonstrations as to how to go about analysing and varying a piece for performance;

(2) providing counter-examples to instantiations of rules, to test whether the student understands limits of application (consider our earlier example of rule conflict);

(3) providing examples which differ in a single respect to test whether the student could actually discern aurally the variations which he is proposing, and for purposes of bringing him to discover his own mistakes aurally;

(4) giving the student an example which deliberately contains 'mistakes' and asking him to 'debug' it;

(5) simulations of common faults on conventional instruments and of musical styles, asking the student to perform the same piece several times in different styles, and finally in altering the musical example (perhaps by making a melody rise instead of fall, all other things being equal) so as to allow him to explore how this affects which rules apply and the musicality of the result (similar to altering voltages across a 'black box' circuit).

Before summing up and reaching some kind of assessment in conclusion, I'd like to briefly identify the three greatest difficulties which my proposals would face (slide 16).

SLIDE 16

SOME PROBLEMS

* WHAT ARE THE RULES?

* HOW FEASIBLE IS AUTOMATED STRUCTURAL ANALYSIS?

* WHAT WILL THE COACH ACTUALLY SAY ?

(i.e HOW CAN WE GIVE EXPLANATIONS AND JUSTIFICATIONS?)

Firstly, it is no easy nor uncontentious matter to determine what the rules are by which music in the Western tradition may be performed. The fact that our coaching system has no intention of actually categorically *prescribing*,and that ultimate recourse for acceptance lies with the musical intuitions which the student brings to the tutorial interaction from listening experience, may ameliorate this problem.

Secondly, computational parsing of tonal music appears to be no less a problem than parsing natural language.

Thirdly,it seems to me to be a special problem as to what our computer would actually *say* when giving explanations and justifications.

A human teacher of performance can rely on his ability and reputation as a performer to persuade the student to heed his advice, especially in the early stages when the student may have few developed musical intuitions and recognised ability to make aesthetic judgements. A student once asked me "but how do you know it should be played like that?" (referring to a piece of renaissance lute music). Often one can find justification in the notated text, but sometimes this appears tenuous (thinking of stylistic features) and one has to ask "but don't you agree it sounds better that way?", and perhaps play parodied examples of apparently alternative readings.

SLIDE 17

CONCLUSIONS

*WHAT'S IN IT FOR MUSIC EDUCATION?

(1) FREEDOM FROM TECHNICAL LIMITATIONS
(2) MAKING MPS EXPLICIT/ FORMALISED -EASIER TO LEARN (ICMC '86 panel)
(3) LEARNING BY DISABLED

*WHAT'S IN IT FOR ITS DESIGN?

(1) GOAL DEFINITION IN OPEN DOMAIN
(2) TUTORING IN OPEN DOMAIN
(3) TESTING MUSIC COGNITION THEORIES
(4) SHEWING THAT COMPUTERS CAN MODEL AREA OF HUMAN AESTHETIC CHOICE AND EMOTIONAL INVOLVEMENT

*PRACTICAL PROSPECTS

(1) GESTURAL CONTROL INTERFACES
(2) SOFTWARE DEVELOPMENT TOOLS
(3) MIDI CONFIGURALBILITY

I'm going to conclude by attempting to assess gains to be made from the proposed research.

For music education, despite the third difficulty outlined above, I would hope to make learning interpretation of music easier and more efficient by providing an explicit formalisation of the subject domain, where the learner was freed from inessential technical limitations.

This special difficulty of learning inexplicit areas of music was drawn attention to by the education panel at this year's International Computer Music Conference. Use of proposed special interfaces could bring the subject area within the reach of many disabled students.

For intelligent tutoring system design, many writers (such as Elsom-Cook) have drawn attention to the fact that no existing tutoring systems offer powerful tutorial guidance in open subject domains. The present research would aim to make a contribution to this problem in the open subject domain of musical performance skills, especially in its method of goal definition by co-operative agreement. In addition it would provide an explicit way of testing theories of music cognition and show that areas of human aesthetic and emotional involvement are tractable to computer-based treatment.

Finally, how practically feasible is this research? Sopisticated gestural and visual interfaces of the kind described have only been developed since the beginning of the 1980's, but by far the largest contribution to making real-time computer music research possible with relatively little hardware investment has been the establishment of the MIDI standard and its stimulation of attendant software development tools. On technological grounds, the proposed research is quite realistically possible.

SLIDE 18

SUMMARY

A GUIDED DISCOVERY ENVIRONMENT IS PROPOSED WHICH WILL ALLOW NOVICE INSTRUMENTALISTS THE FREEDOM TO CREATE THEIR OWN IDEAL PERFORMANCES BY A PROCESS OF COLLECTIVE DECISION BETWEEN COACH AND LEARNER, USING A GRAPHIC EDITING INTERFACE WITH USER-DEFINED COMMANDS. TUTORIAL ADVICE IS GIVEN ON MATCHING OF THIS IDEAL EDIT WITH THE STUDENT'S REAL-TIME PERFORMANCE, USING A GESTURAL CONTROL INTERFACE.

KEY PHRASES

*GUIDED DISCOVERY
* NOVICE INSTRUMENTALISTS
*COLLECTIVE GOAL-DEFINITION
* GRAPHIC EDITING INTERFACE
*USER-DEFINED COMMANDS
*GESTURAL CONTROL INTERFACE
* MATCHING OF REAL-TIME PERFORMANCE AND IDEAL EDIT

Michael Baker
Open University (UK)

December 1986

SOME QUESTIONS, ANSWERS AND SUGGESTIONS

Suggestion 1 "Have you considered the possible link between your ideas and neurological research which suggests that different halves of the brain control analytical and creative thinking, which would correspond to musical analysis and performance?"

Suggestion 2 "When there are two or more apparently acceptable structural analyses (such as in your Mozart example), perhaps the one which is most preferred (by most listeners) can be computed by choosing the analysis which leads to the least 'square cut' or regular grouping. In other words, the one which is most interesting to the listener is the one which is least obvious, but not necessarily 'quirky'".

Question 1 "Is prior musical analysis at all necessary for a performer since there are many performers who are brilliant, but have never thought about what they do or how they do it?"

Answer 1 " My answer comes on two levels: firstly, I should say that if they did *not* perform some sort of analysis, they would risk becoming expert at performing the wrong thing!; secondly, I would make a kind of sideways step and deny that it is even possible to perform without some kind of psychological analysis (in the sense of grouping sensory input and sequences of muscular motor programs in an intelligible way) taking place, and that this analysisis is not of the explicit conscious kind."

Question 2 "Am I right in thinking that you somehow want to teach people to play an instrument by getting the computer to play it for them?"

Answer 2 "In a sense, yes. What I'd like to do is to give novices access to interpretation (which is surely the most interesting and creative part) whilst freeing them from the necessity to have highly developed manual dexterity. Most people vastly underestimate the mental component involved in performance and overestimate importance of visible finger movements - it's well known, for example that mental rehersal of a piece is valuable. Nevertheless, we'd need to retain in a tutoring system those aspects of physical movement which are essential to musical expression, which in practise means careful choice of human-machine interface."

REFERENCES, IN ORDER OF CITATION

Lerdahl, F. & Jackendoff, R. A Generative Theory of Tonal Music (MIT Press,1983)

Sundberg, J. , Fryden, L. & Askenfelt, A. What Tells You the Player is Musical?: An analysis by synthesis of music performance in Sundberg, J. (ed.) Studies of Musical Performance (papers presented at the Royal Swedish Institute of Technology)

Friberg, A. & Sundberg, J. RULLE - in Proceedings of the 1986 International Computer Music Conference

Lamb, M. & Buckley, V. A user interface for teaching piano keyboard techniques Interact '84 ed. Shackel, B. 1985

Freed, A. MacMix - in Proceedings of the 1986 International Computer Music Conference

Mathews, M.V. & Abbott, C. The Sequential Drum Computer Music Journal vol.4, no.4, Winter 1980

Free, J,, Vytas, P. & Buxton, W. What Ever Happened to SSSP? in Proceedings of the 1986 International Computer Music Conference. (see also Pennycook, B. Computer-music Interfaces: A Survey in ACM Computing Surveys vol.17, no.2, June 1985)

Chadabe, J. Interactive Composing: An overview Computer Music Journal 1983

Bamberger, J. Learning to think musically Music Educators Journal, March 1973

Bamberger, J. Music LOGO (Terrapin Inc. Cambridge Massachusetts)

Vercoe, B. The synthetic performer in the context of live performance in Proceedings of the 1984 International Computer Music Conference

CENTRE FOR INFORMATION TECHNLOGY IN EDUCATION

List of CITE Reports

These reports may be obtained from:

Hansa Solanki, Institute of Educational Technology, The Open University, Walton Hall, MILTON KEYNES, MK7 6AA, England.

Report No.	Title and Author
1	A.T. Vincent, (1985) Computing and the Blind.
2	A. Jones, G. Kirkup, J. Morrison (1985) A Trial of Home Based Computer Terminals.
3	Gill Kirkup, (1985) The Present and Potential Use of Ceefax in the Open University.
4	Mark Elsom-Cook, (1986) Artificial Intelligence and Computer Assisted Instruction.
5	Mark Elsom-Cook, (1986) A Pascal program checker.
6	Simon Holland, (1986) How computers are used in the teaching of music and speculations about how Artificial Intelligence could be applied to radically improve the learning of composition skills.
7	Simon Holland, (1986) Design consideration for a human-computer interface using 12-tone three-dimensional harmony space to aid novices to learn aspects of harmony and composition.
8	Alison Petrie-Brown, (1987) The Influence of Context and Coherence as a Foundation for Dialogue Research.
9	Eileen Scanlon, Randall B. Smith (1987) A Rational Reconstruction of a Bubble Chamber Simulation Using The Alternate Reality Kit.
10	Mark Elsom-Cook, (1987) Intelligent Computer-Aided Instruction research at the Open University.
11	Mark Elsom-Cook, (1987) Towards a framework for human-computer discourse.
12	Mark Elsom-Cook, (1987) MATILDA AND IMPART: Lisp tools.
13	Mark Elsom-Cook, (1987) Guided discovery tutoring and bounded user modelling in Intelligent Computer Aided Instruction ed. J. Self, Chapman-Hall 1987.
14	A.M. Petrie-Brown and M.T. Elsom-Cook, (1987) An Examination of an AI model of indirect speech acts.
15	A. Edwards, (1987) Integrating Synthetic Speech With Other Auditory Cues In Graphical Computer Programs For Blind Users.
16	S. Holland. (June 1987) A knowledge-based tutor for music composition.

17 S. Holland, (June 1987) New Cognitive Theories of Harmony Applied To Direct Manipulation Tools for Novices.

18 M. Baker. (July 1987) Intelligent Computer-Aided Instruction and Musical Performance Skills.